

# Using Automated Dependency Analysis To Generate Representation Information

Andrew N. Jackson  
Andrew.Jackson@bl.uk

November 4, 2011

## Abstract

To preserve access to digital content, we must preserve the representation information that captures the intended interpretation of the data. In particular, we must be able to capture performance dependency requirements, i.e. to identify the other resources that are required in order for the intended interpretation to be constructed successfully. Critically, we must identify the digital objects that are only referenced in the source data, but are embedded in the performance, such as fonts. This paper describes a new technique for analysing the dynamic dependencies of digital media, focussing on analysing the process that underlies the performance, rather than parsing and deconstructing the source data. This allows the results of format-specific characterisation tools to be verified independently, and facilitates the generation of representation information for any digital media format, even when no suitable characterisation tool exists.

## 1 Introduction

When attempting to preserve access to digital media, keeping the bitstreams is not sufficient - we must also preserve information on how the bits should be interpreted. This need is widely recognised, and this data is referred to as Representation Information (RI) by the Open Archival Information System (OAIS) reference model [4]. The reference model also recognises that software can provide valuable RI, especially when the source code is included. However, software is not the only dynamic dependency that must be captured in order to preserve access. The interpretation of a digital object may inherit further information from the technical environment as the performance proceeds, such as passwords or licenses for encrypted resources, default colour spaces, page dimensions or other rendering parameters and, critically, other digital objects that the rendering requires. This last case can include linked items that, while only referenced in the original data, are included directly in the performance. In the context of hypertext, the term ‘transclusion’ has been coined to describe this class of included resource [5].

The classic example of a transcluded resource is that of fonts. Many document formats (PDF, DOC, etc.) only reference the fonts that should be used to render the

content via a simple name (e.g. ‘Symbol’), and the confusion and damage that these potentially ambiguous references can cause has been well documented [1]. Indeed, this is precisely why the PDF/A standard [2] requires that all fonts, even the so-called ‘Postscript Standard Fonts’ (e.g. Helvetica, Times, etc.), should be embedded directly in archival documents instead of merely referenced. Similarly, beyond fonts, there are a wide range of local or networked resources that may be transcluded, such as media files and plug-ins displayed in web pages, documents and presentations, or XSD Schema referenced from XML. We must be able to identify these different kinds of transcluded resources, so that we can either include them as explicit RI or embed them directly in the target item (as the PDF/A standard dictates for fonts).

Traditionally, this kind of dependency analysis has been approached using normal characterisation techniques. Software capable of parsing a particular format of interest is written (or re-used and modified) to extract the data that indicates which external dependencies may be required. Clearly, creating this type of software requires a very detailed understanding of the particular data format, and this demands that a significant amount of effort be expended for each format of interest. Worse still, in many cases, direct deconstruction of the bitstream(s) is not sufficient because the intended interpretation deliberately depends on information held only in the wider technical environment, i.e. the reference to the external dependency is implicit and cannot be drawn from the data.

This paper outlines a complementary approach, developed as part of the SCAPE project<sup>1</sup>, which shifts the focus from the data held in the digital file(s) to the process that underlies the performance. Instead of examining the bytes, we use the appropriate rendering software to walk-through or simulate the required performance. During this process we trace certain operating system operations to determine which resources are being used, and use this to build a detailed map of the additional RI required for the performance, including all transcluded resources. Critically, this method does not require a detailed understanding of file format, and so can be used to determine the dependencies of a wide range of media without the significant up-front investment that developing a specialised characterisation tool requires.

## 2 Method

Most modern CPUs can run under at least two operating modes: ‘privileged’ mode and ‘user’ mode. Code running in privileged mode has full access to all resources and devices, whereas code running in user mode has somewhat limited access. This architecture means only highly-trusted code has direct access to sensitive resources, and so attempts to ensure that any badly-written code cannot bring the whole system to a halt, or damage data or devices by misusing them. However, code running in user space must be able to pass requests to devices, e.g. when saving a file to disk, and so a bridge must be built between the user and the protected modes. It is the responsibility of the operating system kernel to manage this divide. To this end, the kernel provides a library of system calls that implement the protected mode actions that the user code needs.

---

<sup>1</sup><http://www.scape-project.eu/>



be instructed to track these child processes, and helpfully reports a brief summary of the command-line arguments that we passed when a new process was launched.

## 2.2 Mac OS X

On OS X (and also Solaris, FreeBSD and some others) we can use the DTrace tool from Sun/Oracle<sup>5</sup>. This is similar in principle to strace, but is capable of tracking any and all function calls during execution (not just system calls at the kernel level). DTrace is a very powerful and complex tool, and configuring it for our purposes would be a fairly time-consuming activity. Fortunately, DTrace comes with a tool called ‘dtruss’, which pre-configures DTrace to provide essentially the same monitoring capability as the strace tool. The OS X kernel calls have slightly different names, the format of the log file is slightly different, and the OS X version of DTrace is not able to log the arguments passed to child processes, but these minor differences do not prevent the dependency analysis from working.

## 2.3 Windows

Windows represents the primary platform for consumption of a wide range of digital media, but unfortunately (despite the maturity of the operating system) it was not possible to find a utility capable of reliably assessing file usage. The ‘SysInternals Suite’<sup>6</sup> has some utilities that can identify which files a process is currently accessing (such as Process Explorer or Handle) and similar utilities (ProcessActivityView, Opened-FilesView) have been published by a third-party called Nirsoft<sup>7</sup>. These proved difficult to invoke as automated processes, and even when this was successful, the results proved unreliable. Each time the process was traced, a slightly different set of files would be reported, and files opened for only brief times did not appear at all. Sometimes, even the source file itself did not appear in the list, proving that important file events were being missed. This behaviour suggests that these programs were rapidly sampling the usage of file resources, rather than monitoring them continuously.

An alternative tool called StraceNT<sup>8</sup> provides a more promising approach, as it can explicitly intercept system calls and so is capable of performing the continuous resource monitoring we need. However, in its current state it is difficult to configure and, critically, only reports the name of the library call, not the values of the arguments. This means that although it can be used to tell if a file was opened, it does not log the file name and so the resources cannot be identified. However, the tool is open source, so might provide a useful basis for future work.

One limited alternative on Windows is to use the Cygwin UNIX-like environment instead of using Windows tools directly. Cygwin comes with its own strace utility, and this has functionality very similar to Linux strace. Unfortunately, this only works for applications built on top of the Cygwin pseudo-kernel (e.g. the Cygwin ImageMagick

---

<sup>5</sup><http://opensolaris.org/os/community/dtrace/>

<sup>6</sup><http://technet.microsoft.com/en-gb/sysinternals/bb842062>

<sup>7</sup><http://www.nirsoft.net/>

<sup>8</sup><https://github.com/ipankajg/ihpublic/>

package). Running Windows software from Cygwin reports nothing useful, as the file system calls are not being handled by the Cygwin pseudo-kernel.

### 3 Results

In this initial investigation, we looked at two example files, covering two different media formats that support transcluded resources: a PDF document and a PowerPoint presentation.

#### 3.1 PDF Font Dependencies

The fonts required to render the PDF test file (the ‘ANSI/NISO Z39.87 - Data Dictionary - Technical Metadata for Digital Still Images’ standards document [3]) were first established by using a commonly available tool, `pdffonts`<sup>9</sup>, which is designed to parse PDF files and look for font dependencies. This indicated that the document used six fonts, one of which was embedded (see Table 1 for details).

The same document was rendered via three different pieces of software, stepping through each page in turn either manually (for Adobe Reader or Apple Preview) or automatically. The automated approach simulated the true rendering process by rendering each page of the PDF to a separate image via the ImageMagick<sup>10</sup> conversion command ‘`convert input.pdf output.jpg`’. This creates a sequence of numbered JPG images called ‘`output-###.jpg`’, one for each page.

All system calls were traced during these rendering processes, and the files that the process opened and read were collated. These lists were then further examined to pick out all of the dependent media files - in this case, fonts. The reconstructed font mappings are shown in Table 1.

The two manual renderings on OS X gave completely identical results, with each font declaration being matched to the appropriate Microsoft TrueType font. The manual rendering via Adobe Reader on Debian was more complex. The process required three font files, but comparing the ‘no-file’ case with the ‘file’ case showed that the first two (DejaVuSans and DejaVuSans-Bold) were involved only in rendering the user interface, and not the document itself. The third file, ‘ZX\_\_\_\_.PFB’, was supplied with the Adobe Reader package and upon inspection was found to be a Type 1 Postscript Multiple Master font called ‘Adobe Sans MM’, which contains all the variants of a typeface that Adobe Reader uses to render standard or missing fonts. Adobe have presumably taken this approach in order to ensure the standard Postscript fonts are rendered consistently across platforms, without depending on any external software packages that are beyond their control.

Although the precise details and naming conventions differed between the platforms, each of the ImageMagick simulated renderings pulled in the essentially the same set of Type 1 PostScript files, which are the open source (GPL-compatible license) versions of the Adobe standard fonts. This is not immediately apparent due to

---

<sup>9</sup>Part of Xpdf: <http://foolabs.com/xpdf/>

<sup>10</sup><http://www.imagemagick.org/>

Tool	Operating System	List of Fonts
pdffonts 3.02	OS X 10.7	Arial-BoldMT, ArialMT, Arial-ItalicMT, Arial-BoldItalicMT, TimesNewRomanPSMT, BBNPHD+SymbolMT (embedded)
Apple Preview 5.5	OS X 10.7	/Library/Fonts/Microsoft/... Arial Bold.ttf, Arial.ttf, Arial Italic.ttf, Arial Bold Italic.ttf, Times New Roman.ttf
Adobe Reader X (10.1.0)	OS X 10.7	/Library/Fonts/Microsoft/... Arial Bold.ttf, Arial.ttf, Arial Italic.ttf, Arial Bold Italic.ttf
Adobe Reader 9.4.2	Debian Linux 6.0.2	/usr/share/fonts/truetype/ttf-dejavu/... DejaVuSans.ttf, DejaVuSans-Bold.ttf /opt/Adobe/Reader9/Resource/Font/ZX____.PFB
ImageMagick 6.7.1	OS X 10.7 via MacPorts	/opt/local/share/ghostscript/9.02/Resource/Font/... NimbusSanL-Bold, NimbusSanL-Regu, NimbusSanL-ReguItal, NimbusSanL-BoldItal, NimbusRomNo9L-Regu
ImageMagick 6.6.0	Debian Linux 6.0.2	/usr/share/fonts/type1/gsfonts/... n019004l.pfb, n019003l.pfb, n019023l.pfb, n019024l.pfb, n021003l.pfb
ImageMagick 6.4.0	Cygwin on WinXP	/usr/share/ghostscript/fonts/... n019004l.pfb, n019003l.pfb, n019023l.pfb, n019024l.pfb, n021003l.pfb

Table 1: Font dependencies of a specific PDF document, as determined via a range of tools.

the different naming conventions using on different installations, but manual inspection quickly determined that, for example, NimbusSanL-Bold and n019004l.pfb were essentially the same font, but from different versions of the gsfonts package. The information in the system trace log made it easy to determine how ImageMagick was invoking GhostScript, and to track down the font mapping tables that GhostScript was using to map the PDF font names into the available fonts.

Interestingly, as well as revealing that these apparently identical performances depend on different versions of different files in two different formats (TrueType or Type 1 Postscript fonts), the results also show that while Apple Preview and ImageMagick indicate that Times New Roman is a required font (in agreement with the pdffonts results) this font is not actually brought in during the Adobe Reader rendering processes. A detailed examination of the source document revealed that while Times New Roman is declared as a font dependency on one page of the document, this appears to be an artefact inherited from an older version of the document, as none of the text displayed on the page is actually rendered in that font.

### 3.2 PowerPoint with Linked Media

A simple PowerPoint presentation was created in Microsoft PowerPoint for Mac 2011 (version 14.1.2), containing some text and a single image. When placing the image, PowerPoint was instructed to only refer to the external file, and not embed it, simulating the default behaviour when including large media files. The rendering process was then

performed manually, looking through the presentation while tracing the system calls. As well as picking up all the font dependencies, the fact that the image was being loaded from an external location could also be detected easily.

The presentation was then closed, and the referenced image was deleted. When re-opening the presentation, the system call trace revealed that PowerPoint was hunting for the missing file, guessing a number of locations based on the original absolute pathname. This approach can therefore be used to spot missing media referenced by PowerPoint presentations.

## 4 Conclusions

Process monitoring and system call tracing is a valuable analysis technique, complementary to the more usual format-oriented approach. It enables us to perform detailed quality assurance of existing characterisation tools, using a completely independent approach to validate the identification of the resources required to render a digital object. Furthermore, because the tracing process depends only on standard system functionality, and not on the particular software in question, it can work for all types of digital media without developing software for each format. As the PowerPoint example shows, the only requirement for performing this analysis is the provision of suitable rendering software.

Before using this approach in a production setting, it will be necessary to test it over a wider range of documents and types of transclusion, e.g. embedded XML Schema. In particular, the monitoring should be extended to track network requests for resources as well as local file or software calls. Although all network activity is visible via kernel system calls, the raw socket data is at such a low level that it is extremely difficult to analyse. Fortunately, tools like netstat<sup>11</sup> and WireShark<sup>12</sup> have been designed to solve precisely this problem, and could be deployed alongside system call tracing to supply the necessary intelligence on network protocols. Beyond widening the range of resources, extending this approach to the Windows platform would be highly desirable. The current lack of a suitable call tracing tool is quite unfortunate, and means that this approach cannot be applied to software that only runs on Windows. Hopefully, StraceNT can provide a way forward.

Beyond the direct resource dependencies outlined here, this approach could be combined with knowledge of the platform package management system in order to build an even richer model of the representation information network a digital object requires. For example, Debian has a rigorous package management processes, and by looking up which packages provide the files implicated in the rendering, we can validate not only the required binary software packages, but also determine the location of the underlying open source software, and even the identities of the developers and other individuals involved. This allows very rich RI to be generated in an automated fashion. Furthermore, as the Debian package management infrastructure also tracks the development and discontinuation of the various software packages, this information could be leveraged to help build a semi-automatic preservation watch system.

---

<sup>11</sup><http://en.wikipedia.org/wiki/Netstat>

<sup>12</sup><http://www.wireshark.org/>

## 5 Acknowledgments

This work was partially supported by the SCAPE Project. The SCAPE project is co-funded by the European Union under FP7 ICT-2009.4.1 (Grant Agreement number 270137).

## References

- [1] G. Brown and K. Woods. Born Broken : Fonts and Information Loss in Legacy Digital Documents. *International Journal of Digital Curation*, 6(1):5–19, 2011.
- [2] International Standardization Organization. ISO 19005-1:2005 Document management – Electronic document file format for long-term preservation – Part 1: Use of PDF 1.4 (PDF/A-1), 2005.
- [3] National Information Standards Organization. ANSI/NISO Z39.87 - Data Dictionary - Technical Metadata for Digital Still Images, 2006.
- [4] The Consultative Committee for Space Data. Reference Model For An Open Archival Information System (OAIS), 2009.
- [5] Theodor Holm Nelson and Robert Adamson Smith. Back to the Future, 2007.